Expressive Visualization and Rapid Interpretation of Seismic Volumes

Daniel Patel

Thesis for the degree of Philosophiae Doctor (PhD)
at the University of Bergen, Norway

August 2009

Acknowledgement

I would like to thank the people who supported me during my PhD:

To my mother and father for unlimited love. To my brother for guidance in both work and in life, for leading the way and demonstrating where not to go. To my aunt, uncle and family in Connecticut for unreserved hospitality and for offering me and my friends a peaceful place for rest and care.

To Ivan for being a fine friend and for advice and inspiration in research. To Kha for long lasting friendship, Bård for deep conversations, Eirik for supporting my decision, Silja for caring, Magnus for great fun, Katarina for sophistication, Kristina for sharing thoughts, Christian for always being there, Mathilde for being sweet, Egle for good memories, Inga-Karin for being a friend, Ulf for physical development and good conversations and Ingelin for adventures. My thanks also go to Raymond, Agathe, Stefan in Frankfurt, Halvard, Gergø and Hege.

Scientific environment

This PhD has been made possible through a scholarship from Christian Michelsen Research. I am grateful to CEO Arvid Nøttvedt and vice presidents Kåre Villanger and Christopher Giertsen for standing behind this initiative. They identified one of the leading visualization research groups in Europe; the visualization group at the Technical University of Vienna, where I spent about half the PhD under my thesis supervisor Meister Eduard Gröller. Thanks Meister and CMR for excellent supervision and for giving me freedom to develop my own ideas.

The thesis topic of illustrative visualization was directly inspired by the work of Stefan Bruckner, Ivan Viola and Peter Rautek in the Vienna group. Much knowledge and fun came from them and the other group members: Erald Vucini, Peter Kohlmann, Martin Haidacher who coauthored a paper, Muhammed Muddasir Malik, Maurice Termeer and Matej Mlejnek. The fine VolumeShop software developed by Stefan Bruckner has mainly been the development framework for the techniques presented in this thesis.

The other half of my PhD was spent at the then newly created group at the University of Bergen now led by Professor Helwig Hauser and Ass. Prof. Ivan Viola. I had several good discussions with the professors and the PhD students Johannes Kehrer, Ove Daae Lampe who is also sponsored by CMR and Paolo Angelelli. The PhD student Jean-Paul Balabanian who coauthored a paper helped me a lot with ideas and support. The work by the Master student and coauthor Øyvind Sture has also been very valuable in my PhD.

Without guidance from domain experts in the oil and gas industry my research would not be possible. These people work on very tight schedules, so thanks to my co-supervisor Dr. Christopher Giertsen for identifying and getting me in contact with the following experts in StatoilHydro who were willing to discuss new ideas: Thanks to Dr. John Thurmond for listening to ideas and coming back with real scenarios and data to demonstrate them on. Thanks to the geologist and illustrator John Gjelberg for giving me a glimpse into the art of illustration drawing and for showing how a 'geoillustrator' does his work. Thanks to Jens Grimsgaard for supporting the project.

List of publications

Illustrative Rendering of Seismic Data
Published in: Vision, Modelling and Visualization (VMV) 2007
Authors: Daniel Patel, Christopher Giertsen, John Thurmond, Eduard Gröller

The Seismic Analyzer: Interpreting and Illustrating 2D Seismic Data
Published in: IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 6, November/December 2008
Authors: Daniel Patel, Christopher Giertsen, John Thurmond, John Gjelberg, and Eduard Gröller

Moment Curves
Published in: Proceedings of Pacific Vis, April 2009
Authors: Daniel Patel, Martin Haidacher, Jean-Paul Balabanian, Eduard Gröller

Seismic Volume Visualization for Horizon Extraction
Technical report/manuscript, 2009
Authors: Daniel Patel, Stefan Bruckner, Ivan Viola, Eduard Gröller

Knowledge-assisted visualization of seismic data
To be published in: Computers and Graphics Journal
vol 33 issue 4, scheduled to appear end-September / beginning-October 2009
Authors: Daniel Patel, Øyvind Sture, Helwig Hauser, Christopher Giertsen, Eduard Gröller

List of presentations

Paper I presentation: Saarbrücken, Germany, November 7, 2007

Knowledge-assisted Visualization Workshop 'Knowledge-assisted visualization of seismic data', IEEE VisWeek, Colombus, Ohio, USA, October 19, 2008

Paper II presentation: IEEE VisWeek, Colombus, Ohio, USA, October 23, 2008

Arranger/chair on half day session 'State of the art Visualization in Geosciences'.
European Geosciences Union, General Assembly 2009, Vienna April 20, 2009
(Self)invited talk: 'Illustrative Visualization for Rapid Interpretation in Geosciences'

Paper III presentation: Pacific Vis, Peking University, Beijing, China, April 23, 2009

Invited talk: IllustraVis Workshop, University of Bergen 'Illustrative Visualization for Rapid Interpretation in Geosciences', June 4, 2009

Invited talk: VRGeo Consortium, Fraunhofer IAIS, 'Illustrative Visualization of Seismic Data', Sankt Augustin, Germany, June 16, 2009

Contents

Chapter 1

Abstract

One of the most important resources in the world today is energy. Oil and gas provide two thirds of the world energy consumption, making the world completely dependent on it. Locating and recovering the remaining oil and gas reserves will be of high focus in society until competitive energy sources are found. The search for hydrocarbons is broadly speaking the topic of this thesis. Seismic measurements of the subsurface are collected to discover oil and gas trapped in the ground. Identifying oil and gas in the seismic measurements requires visualization and interpretation. Visualization is needed to present the data for further analysis. Interpretation is performed to identify important structures. Visualization is again required for presenting these structures to the user. This thesis investigates how computer assistance in producing high-quality visualizations and in interpretation can result in *expressive visualization and rapid interpretation of seismic volumes*. Expressive visualizations represent the seismic data in an easily digestible, intuitive and pedagogic form. This enables rapid interpretation which accelerates the finding of important structures.

Chapter 2

Introduction

The diverse work of the five articles constituting this thesis has been divided into the categories of expressive visualization and rapid interpretation. Expressive visualization is one of the two enabling technologies for rapid interpretation. The other enabling technology is automatic structure identification algorithms. Such algorithms segment data into meaningful objects. Certain expressive visualizations makes use of segmented data by applying different presentation styles on each object in the data. In the geoscientific domain, segmentations are typically obtained through an interpretation process. Thus it is not a strict ordering with visualization first and then interpretation. Visualization and interpretation goes in a spiral. Visualization enables interpretation, which leads to more information about the data that can be used to produce better visualizations. This in return gives more insight into the data.

2.1 Expressive Visualization

Expressive visualizations present data in a more expressive and clearer way than standard visualization methods such as slice rendering, polygonal surface rendering and volume rendering do. Expressive visualizations are perceptually more easy to digest. This is achieved by shaping the presentation of the data to how the human perceptual system works. Previous studies and knowledge from art, illustration techniques and perceptual psychology has been used to develop methods which efficiently present data in a communication friendly form. We are applying these methods to create expressive visualizations. Illustrations are examples of expressive visualizations. Illustrations are frequently used for communicating completed seismic interpretations. The ability of illustrations to express abstractions make them well suited for representing the high-level knowledge that is produced during interpretation. In the oil and gas industry, high-quality illustrations have a central role in the final communication of interpretation results for decision making and in reporting. The expressive visualizations presented in this thesis attempt to:

- reduce the time to create geoscientific illustrations.
- make it easier for interpreters to express their hypotheses and inner models.
- offer new visualization methods tailored for seismic data.

The three bullet points are in order discussed in the next three sections.

Computer-generated Illustrations Geoscientific illustrations are ideal for communicating interpretations to decision makers and to the public. They are extensively used in reports when bidding for concessions to perform surveys and drills. The illustrations are also important in educational material for geosciences.

Automated and tailored tools for creating geoscientific illustrations are lacking. The creation process is time consuming and manual. Manually drawing illustrations binds valuable human resources in an oil company as it requires special competences from geologists with illustration knowledge. We attempt to increase the efficiency of geoscientific illustrators by automating common and time consuming drawing tasks such as sketching textures and drawing suggestive horizons. These tasks are currently performed rather manually in general purpose drawing software. In contrast, automating this procedure results in the generation of illustrations in fractions of a second after a model is defined. The sub-second rendering time gives rise to a new mode of interactivity in illustrations. Tasks such as changing the angle of the illustration, the perspective, the texture types or the cut planes takes place with immediate results. This is an important improvement from manual illustrations where performing such changes requires drawing a new illustration and may take hours or days. Interactivity in illustrations increases their communicative impact compared to classical static illustrations.

Hypothesis Externalization Instead of only applying the illustrative techniques on finished interpretations, we have considered the effects of using computer assisted illustrative sketching techniques *during* interpretation. The quality of seismic data can vary, the data can be noisy and unclear. Consequently the interpreters must compose hypotheses of possible scenarios and make educated guesses. Interpretation is regularly performed in multidisciplinary teams and it is important that the members understand each others ideas. By allowing quick sketching of illustrations through computer assistance, the interpreters are given a tool for expressing and sharing their internal hypotheses and models. Externalizing hypotheses provides the means for discussing ideas and arriving at a common ground of understanding.

Improved Seismic Visualization The visualization research communities have had little focus on seismic data. Promising results have been achieved within subjects such as medical visualization and information visualization. We attempt to reduce this gap by adapting and tailoring methods from other visualization fields and developing new visualization methods tailored for seismic data such as:

- improved 3D seismic visualization using real-time approximated ambient occlusion volume rendering.
- 3D sparse texturing for conveying rock layering and rock types in true 3D.

2.2 Rapid Interpretation

Interpretation is required when searching for oil and gas. Oil and gas is created when organic material is deposited, buried and pressurized under high temperatures over long periods of time. The oil and gas will then migrate upwards and may get trapped in tightly sealed reservoir structures and accumulate there. For finding the oil or gas, signs of the producing deposits or collecting seals are searched for with interpretation techniques. The aspects of rapid interpretation in this thesis attempt to reduce the manual interpretation work and increase interpretation speed by:

- providing overviews of the data so interpretation in unimportant areas is avoided and focus can be made on areas with potential.
- reducing time spent on repetitive and time consuming tasks by automating them so the interpreter can concentrate on challenging areas in the seismic data where automatic methods fail.
- making it quicker and easier to verify a finished interpretation by comparing it and visualizing it aligned and in context with the underlying data the interpretation was based on.

These points of rapid interpretation are further elaborated in the three following sections.

Overviews One element enabling rapid interpretation is the ability to gain an overview of the data. Interpretation is typically time consuming due to focus on high accuracy at the start. The reason for high accuracy at the very start stems partly from the limited possibilities to create overviews in this early stage. Overviews are difficult to obtain because of the large seismic volumes collected from extensive land or sea areas. Several physical measurements for the same area also cause obtaining overviews challenging. Overviews are accomplished in this thesis with methods for compact visualization of large areas and with methods for the simultaneous visualization of several measurements in overlapping areas. Large areas are visualized by constructing sparse representations containing only the absolutely necessary information and presenting this information in zoomed-out thumbnail views. Overlapping volumetric measurements are presented by mapping the volumes to disparate representations that are perceptually separable, followed by merging the representations carefully into one view. This procedure provides the means for multiattribute visualizations.

Automatic Structure Identification Another element enabling rapid interpretation is having automatically extracted structure suggestions at hand before manual interpretation begins. This allows the interpreter to quickly browse through the suggestions and select good ones. In contrast, in the current workflow the interpreter is either manually identifying and marking structures or, when working in a semiautomatic mode, the interpreter is constantly interrupted by the initialization of structure identification algorithms followed by waiting for the automatic processing to finish.

Verification Interactivity in illustrations allows easier verification of the interpretation represented in the illustration. In an interactive rendering the user can switch from seeing the original uninterpreted data to seeing the interpreted illustration. Smoothly fading out the original data while fading in the illustration makes it possible to compare the underlying data with the interpretation and this enables verification.

Chapter 3

Results

This thesis consists of 4 published articles and one unpublished technical report. Each article emphasizes parts of the elements described in the introduction. Using the same order and structure as the introduction, this chapter presents the results we have achieved. References to the articles for further reading is given.

3.1 Results in Expressive Visualization

The techniques for constructing expressive visualizations have been developed through the study of subjects such as art, illustration techniques and human perception. The following works give a broader overview of the underlying literature and science behind creating expressive visualizations. 'Semiology of graphics' [5] written by Bertin in 1967 can be considered a classic work laying the foundation of information visualization. Bertin studied the visual symbols used in drawings and graphics such as lines, patterns, stipplings and colors. He showed how these symbols could be combined to achieve information-rich visualizations. Tufte wrote the book 'The Visual Display of Quantitative Information' [14] about efficient presentation of information. He takes numerous examples of good and bad charts and diagrams from newspapers and analyses them. Guidelines for efficient visual communication of information is presented in his works. A scientific approach towards perception is taken by Margaret S. Livingstone in the book 'Vision and Art: The Biology of Seeing' [13]. Livingstone is a neurobiologist, she explains why the brain more easily understands certain visual representations over others and makes connections to techniques used in art.

Before we present our results in expressive visualizations we start with describing the seismic data and how it can be visualized using basic techniques. The data we are visualizing in this thesis is mainly seismic reflection data. This data is acquired by sending sound waves into the earth and processing the reflected sound. This is analogous to how ultrasound measurements are performed in medicine. A basic seismic survey results in a 2D image of a vertical slice into the ground. More advanced 3D surveys result in a series of 2D vertically stacked slices spanning a 3D reflection volume of the earth.

From the seismic reflection data, new data can be calculated. The new data will emphasize a certain property or *attribute* of the original data. Since the new data is *derived* from the original, such data is called *derived attributes.* Further information about seismic attributes can be found in the work by Iske and Randen et al. [11].

Visualizing data requires a mapping from the numerical data values to the 2D grid of colored pixels constituting the computer display. The visualization of a 2D seismic slice is created by mapping from the measured slice values to colors. Several color mappings can be used. For seismic reflection data it is common to map the range of low-to-high reflection values to a color transition from red via white to blue or from black to white, see Figure 1. The look-up table that defines the mapping from data values to a visual presentation is called a transfer function. See Figure 2 for a 3D seismic reflection dataset rendered with a red-white-blue transfer function and Figure 3 top left where a white-black transfer function is used. The transfer function can also assign degrees of transparency to the value range. Transparencies result in blending and visualization of the data values behind the transparent ones, this creates a 3D rendering of the volume. Hadwiger et al. [10] give an in depth introduction to the principles of volume rendering. To enable the appearance of objects in 3D space as if the computer screen is a look-through window into a 3D world, translations and rotations are performed on the data followed by a projection to the 2D computer screen. The reader is referred to Foley et al. [8] for basic concepts of transformations and the rendering pipeline.

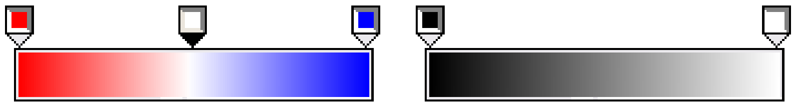

Figure 1: Red-white-blue and black-white transfer functions

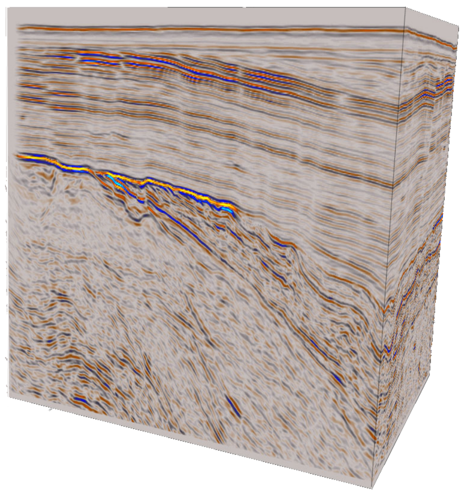

Figure 2: Red-white-blue seismic reflection volume. No reflection values have been set to transparent, therefore the volume can not be looked into and only the side surfaces of the volume block is visible.

Basic visualization techniques were briefly discussed in this section. The following section will present more advanced visualization techniques conceived in this thesis.

3.1.1 Computer-generated illustrations

The drawing techniques for creating illustrations have been developed for the purpose of communicating high-level aspects in a simple way. Elements of interest are illustratively emphasized to represent the knowledge acquired during some analysis process. Illustrative methods allow for annotating the data so that the resulting image communicates its intent clearly to the viewer. Figure 3 gives a comparison of computer-generated raw data visualization in a) vs a manually crafted illustration in b). Figure 3c) is a computer-generated result from this thesis' work.

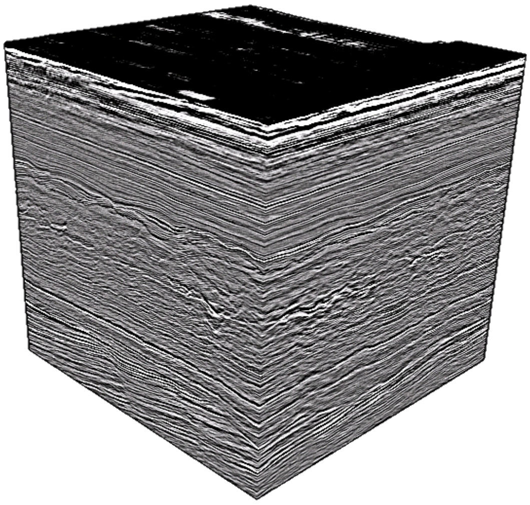

(a) Raw seismic data

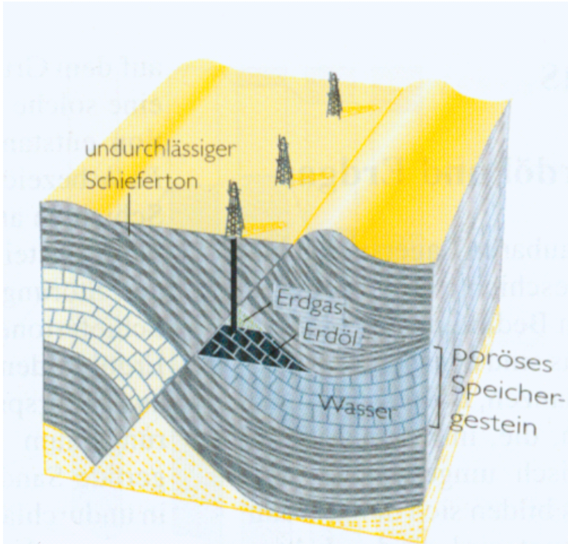

(b) Manually drawn illustration

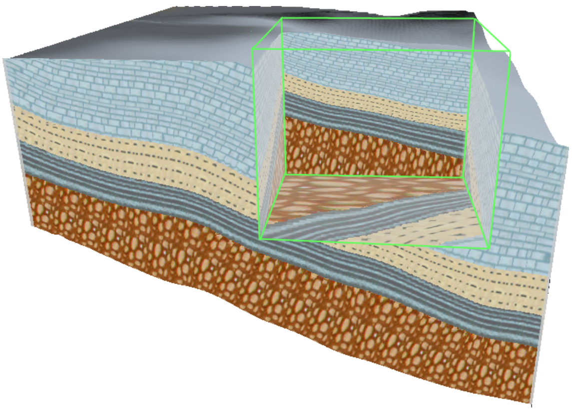

(c) Computer generated illustration

Figure 3: a) Raw data. b) Manual illustration from 'Understanding Earth' [9] showing an oil and gas seal made from a faulted impermeable rock layer. c) Our computer-generated illustration where the content of the green wireframe box has been removed to enable insight into the volume.

As can be seen in the handcrafted illustration taken from geoscientific literature, textures are used extensively. There are several reasons for this. As opposed to color-coding, textures are perceptually more expressive. Textures communicate the rock type and the rock orientation. Compression can be presented by varying the density of the texture throughout the layer. Erosion and faults can be shown using discontinuities in the textures (a fault is shown as the diagonal crack in Figure 3b). Information is embedded directly in the texture. In addition, textures give aesthetically pleasing results. Reading out information without having textures requires more work. Imagine Figure 3b) without textures thus only having line borders separating the layers and the fault. Identifying the rock type would now require searching through the layer for a text label or for a line going to a text label describing the rock type. Even if the layers would be color-coded the user would still have to look up the rock type in a color legend. The angle of the rock at any position would now be deduced by finding and averaging the angle of the upper and lower boundary line of the layer close to that position. Compression would be deduced by looking at the varying distance of the upper and lower boundary line. Thus without textures, the information is no longer locally presented but must be searched for in the image. For interactive visualizations further problems arise when trying to deduce layer angles and compression if the layer boundaries are not visible on screen due to high zoom levels.

These are some reasons why the use of textures was adopted and integrated into the geosciences. To ease communication, geologists and geoillustrators use a standardized language for expressing their knowledge. This language consists of different textures for representing information such as rock types. The US Federal Geographic Data Committee has produced a document [2] with over one hundred of such standardized textures. Figure 4 shows three pages from this geological texture library. The patterns are cleverly designed

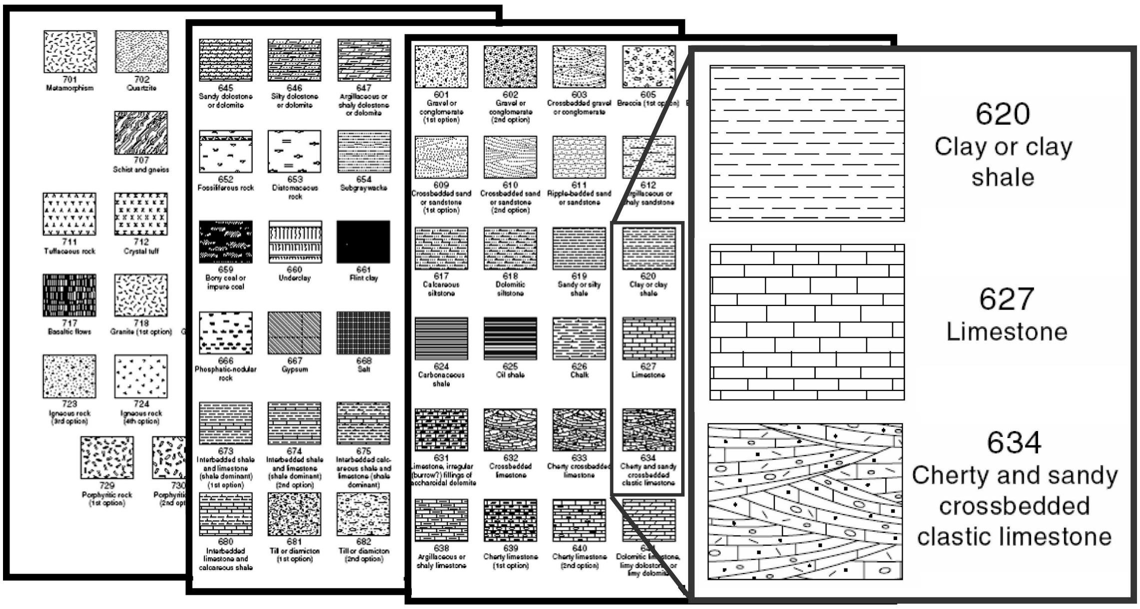

Figure 4: Common symbology in geosciences. Rock textures [2].

and constitute a taxonomy. Rock types of the same subgroup are represented with similar patterns. Thus, even if a geologist does not recognize the exact pattern she is looking at, she should be able to identify the general group it belongs to when knowing other rocks with similar patterns.

On earth, rock layers are laid down in succession and they represent successive geological time periods. A widely used and standardized table of color-codings representing geologic time from the current Quaternary period to the Archean eon 2.5 billion years ago is also part of the 'visual' or 'symbolic' language used by geoscientists (see Figure 5). This language has been developed by the Commission for the geological map of the world [1].

In contrast to other domains where advanced visualization is frequently used such as in medicine, geoscientists heavily make use of a standardized visual language for communication. Therefore expressive visualizations such as illustrations have a large potential in the geoscientific domain. One of the goals of this thesis has been to automate the techniques required for creating geoscientific illustrations and integrate the use of illustrations

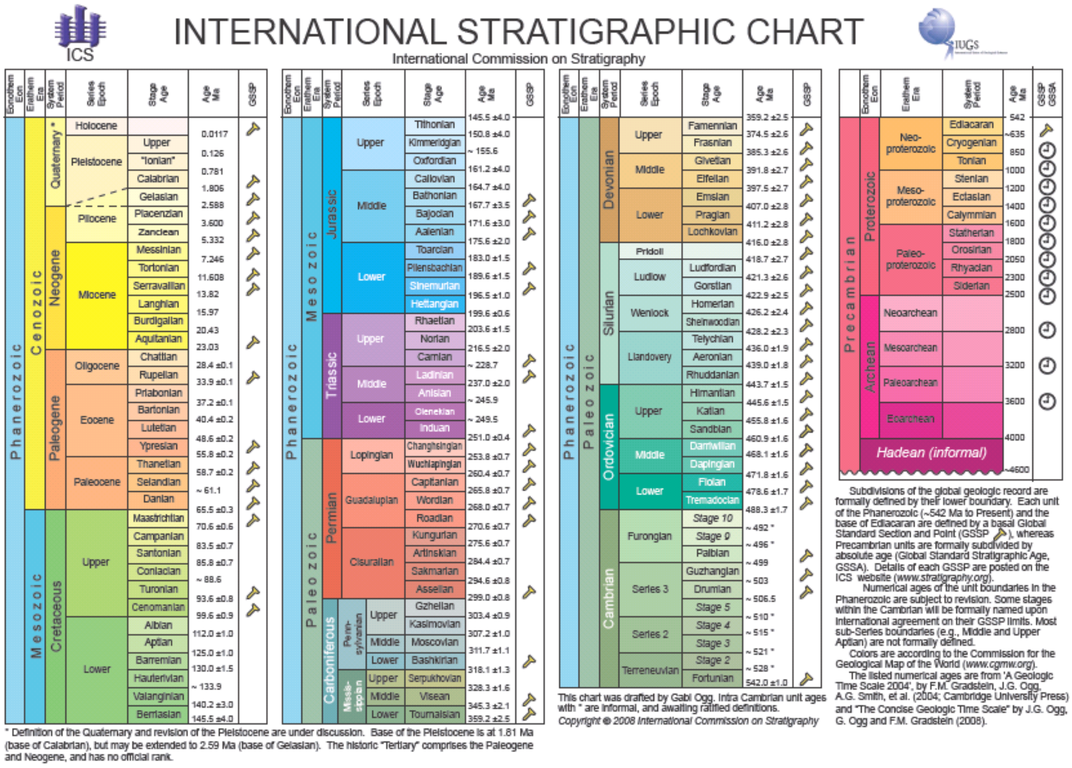

Figure 5: Common symbology in geosciences. Geologic time colors [1].

in the work flow. Achieving the expressive visualizations that these illustrations represent, requires two components. One must assign textures to the data and one must specify how the textures will bend. These two components are elaborated on in the following two sections.

Texture Transfer Functions In our work we achieve the texturing seen in illustrations by using transfer functions extended to handle textures. We call these 'texture transfer functions'. Similar to color transfer functions which assign colors to attribute values, texture transfer functions assign textures to attribute values. Figures 6a) and b) show examples of two texture transfer functions to the left and the respective results when applied on the seismic data to the right. The textures shown in squares above the transfer functions to the left blend into each other analogous to how the colors in Figure 1 shown above the transfer functions blend into each other. In the simplified examples of Figure 6, the horizontal axis of the transfer function is not mapped to a seismic attribute, but is simply mapped to the horizontal axis of the result image to the right. Therefore textures to the left in the texture transfer function are seen to the left of the result image and textures to the right are seen to the right in the result image. In addition to assigning textures, transparencies are assigned along the horizontal axis of the transfer function. The transparencies are defined by the height of the blue curve in the transfer function. When the curve is at its vertically lowest in Figure 6 left, the texturing is completely transparent and when the curve is at its vertically highest, the texturing is completely opaque. By using transparencies it becomes possible to look through the texture and directly at the data displayed behind. By varying how textures and transparencies are assigned to the data values, several effects can be obtained. Figure 6a) shows a transfer function that softly fades in from complete transparency where the underlying seismic data is visible, to complete opaqueness and then blends between textures. Textures bend according to the underlying seismic data displayed above the result images. To achieve this effect the data must be analyzed so the 'bending' can be extracted and applied to the textures. Extraction of the bending information is described in the next section. In Figure 6b) an abrupt change from transparent to opaque takes place due to the staircase shape of the opacity curve in the transfer function. The opaque textures then blend from few random lines to textures with an increasing number of lines. In example b) the textures do not bend according to the underlying seismic as this is not appropriate for all types of textures.

Figure 7 shows a slice through a dataset manually segmented into four rock layers. The layer segmentation is the same as shown in Figure 3c) although other textures are used. Each rock layer is assigned a unique texture. Two textures controlled by an attribute is shown on top of the rock layer textures. They are a boulder texture (blue) for low values and a brick texture (yellow/pink) for high values of the attribute. The high valued interval represented with the brick texture is further differentiated by blending in a yellow-to-pink color transfer function. This distinguishes between areas of low values in yellow, to high

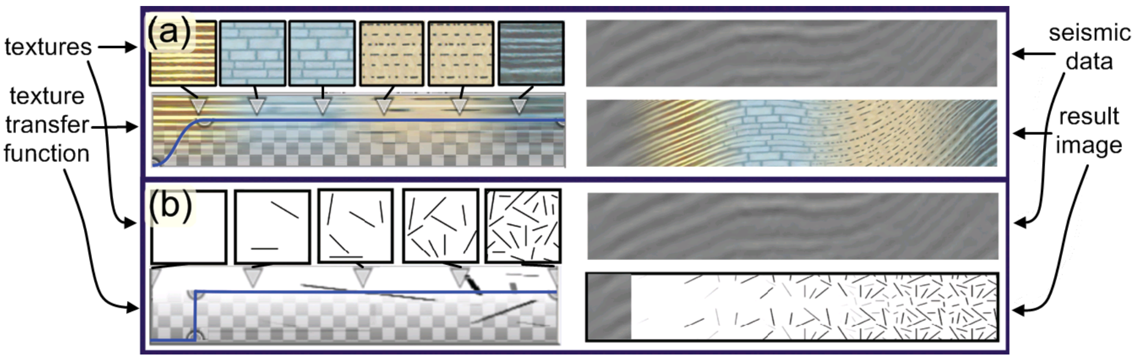

Figure 6: Two texture transfer functions left in a) and b) and the results to the right. The slices in the top right corner of a) and b) show the underlying seismic data.

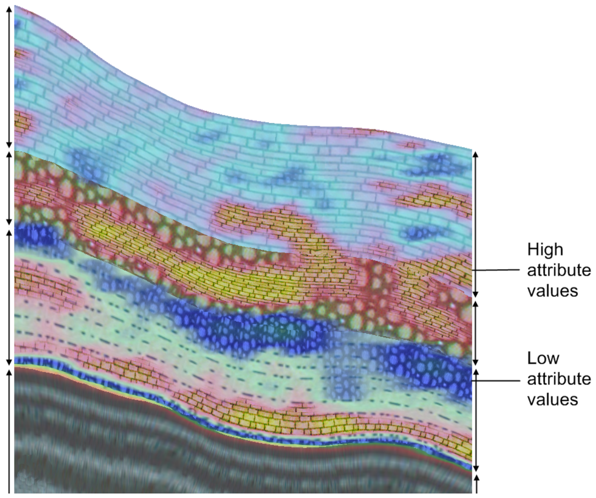

Figure 7: Color and texture transfer functions combined. Arrows denote layer extents.

values in pink for the value interval. Thus the high valued interval is expressed with a specific texture and the value range in this interval is shown with color. Color and texture can be visually mixed and then perceptually separated due to their perceptual orthogonality. The example demonstrates how the use of color and texture transfer functions can control the amount of information that needs to be shown in different parts of attribute value ranges. However the resulting visualization suffers from visual overload and it is difficult to see the underlying rock layering.

Figure 8a) shows another approach for combining colors and textures. A slice with three rock layers is shown. In the middle layer, four intervals of a derived attribute is presented. The derived attribute is shown in Figure 8b) with a color transfer function dividing the attribute values into low (blue), middle (yellow), high (red) and highest (green) values. Mapping the four intervals to different textures would hide the texturing that communicates the rock type for each layer as was shown in the previous example. To keep the rock texture for each layer, all intervals in the middle layer are mapped to the same brick texture in Figure 8a). The intervals are separated by having differently sized textures. The sizes are chosen so the four intervals are with increasing values mapped to increasingly smaller brick textures. The result is brick sizes giving an intuitive visual ordering of the four intervals while at the same time not hiding the layer texturing. The color transfer function used in Figure 8b) is blended in to further contrast the intervals. As can be seen in this example, texture type and texture size can be considered perceptually orthogonal. More examples of texturing for multiattribute visualization will be presented

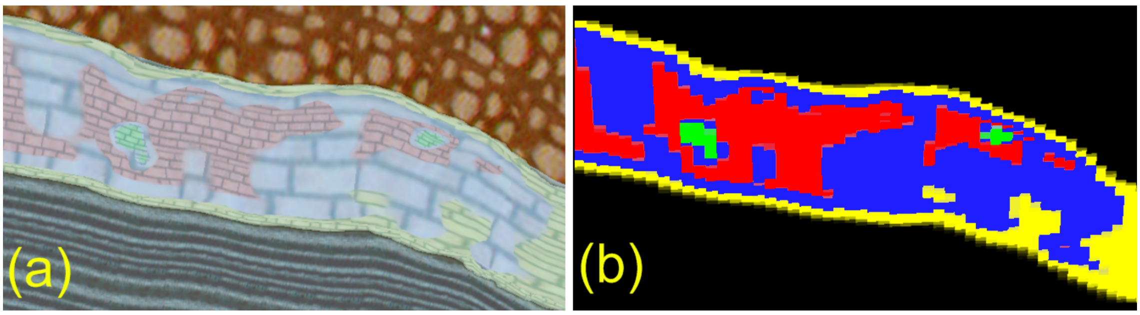

Figure 8: Color and texture transfer functions combined with varying texture density.

in section 3.2.1.

Parameterization To achieve bending textures as found in illustrations, we need to extract bending information from the seismic data, store the information in an appropriate representation, and use the information to deform the textures. The information describing the bending is encoded in a parameterization. One way to represent the parameterization

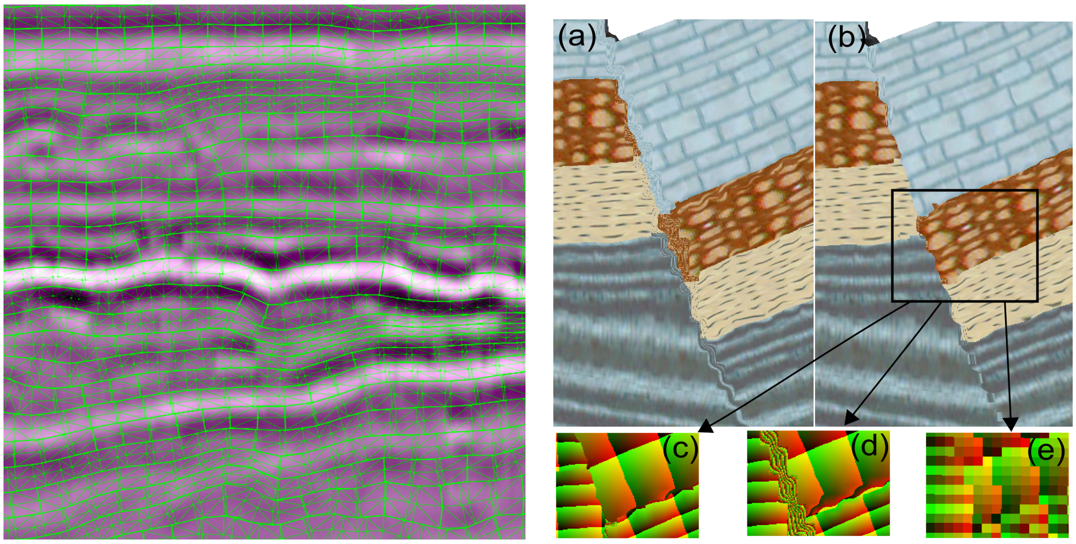

Figure 9: Comparison of parameterization methods of Paper I to the right and Paper II to the left represented as the green overlaid grid.

is by using a grid that is warped to the trend of the seismic data as seen in green in Figure 9 left. The grid spans a deformed Cartesian coordinate system. Drawing textures in this coordinate system will achieve the wished texture bending. A stylized example of Figure 9 left is shown in Figure 10a). Two red lines exemplify a structure in the reflection data and a deformed grid is shown that follows the bending and contraction of the structure. The deformed grid of 10a) is undeformed in Figure 10c). Drawing an undeformed texture in this undeformed Cartesian coordinate system is trivial. With the pair of deformed and undeformed grids, a mapping is created that defines the parameterization. This is the representation that was used in Paper II.

A parameterization representation inverse to the one just described was used in Paper I. Here the parameterization is stored in a regular grid. Basically, for each position of a volumetric reflection value, an additional value is stored describing which coordinate from the texture to use there. In Figure 9b) can be seen a texturing and in 9e) is seen a regular grid of colored pixels representing texture coordinates . This is exemplified in Figure 10b) where a regular grid is shown and in 10d) where the texture lookup coordinates for the corresponding grid points are shown. The approach has weaknesses. A problem arises when trying to represent discontinuous texture mappings such as over the fault in Figure 9b). Linearly interpolated texture coordinates stored in a regular grid disallows discontinuous texture mappings. The discontinuous texture mapping over the fault must be represented by a discrete value jump in the texture coordinates. Due to linear interpolation, a continuous value interpolation between unrelated texture coordinates takes place instead. The result of this erroneous linear interpolation of 9e) is shown in 9d). The method makes it impossible to represent discontinuous texture mappings which are needed when texturing

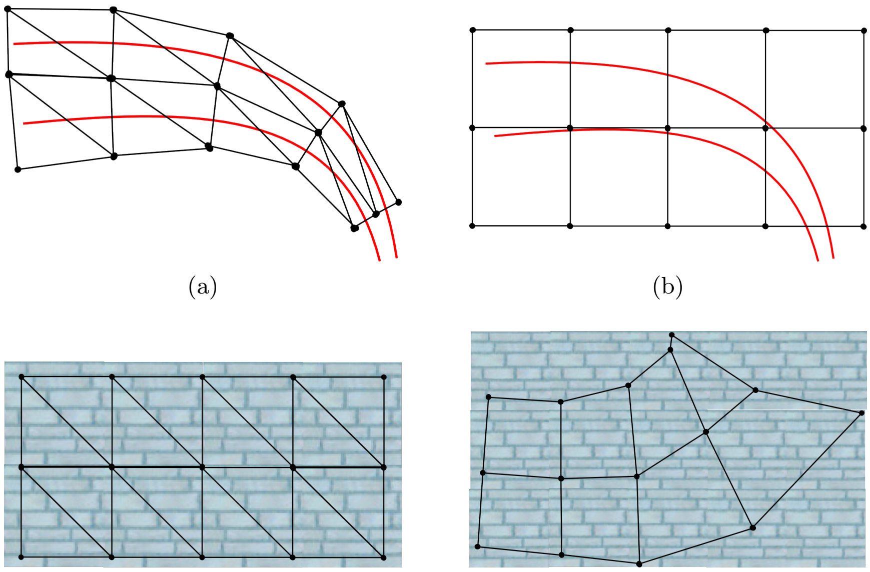

(a) (b) (c) (d)

Figure 10: Conceptualized comparison of parameterization methods of Paper I in a) and c) and Paper II to in b) and d).

over faults. The textural artifacts created over a fault with this method are shown in 9a). In 9b) and 9c) our nonoptimal attempt to solve these problems is shown. Further information of this parameterization can be found in Paper I. Since the solution was not optimal, the better parameterization representation first described was developed for the follow up paper (Paper II).

Two possible representations of the parameterization were discussed above. There are also several ways to extract the parameterization from the underlying data. Structures spanning the data is required for dictating the parameterization. In our works these structure have either been manually interpreted horizons or horizons automatically extracted from the uninterpreted reflection data. Paper I uses manually interpreted horizon surfaces to create a 3D parameterization. In Paper II we create a 2D parameterization using horizons lines automatically extracted from uninterpreted 2D data. A natural next step would be creating a 3D parameterization from uninterpreted 3D data. We believe that the distance transform from the extracted 3D horizon patches used in Paper IV can be a promising starting point for this.

3.1.2 Hypothesis Externalization

Expressive visualization techniques can be used to create geoscientific illustrations that convey finished interpretations. Creating illustrations from completed interpretations has been the current work mode for geoscientific illustrators. We have investigated the effects of illustration sketching early in the workflow, while interpretation is being performed, instead of after it is finished. Such an approach can be useful for communicating ideas and hypotheses during interpretation. This can be an advantage as compared to only having expressive visualizations available after both the interpretation is finished and the manual illustration has been made. An example of a hypothesis quickly sketched during interpretation is shown in Figure 11. Different textures have semitransparently been overlaid on a seismic slice to express one possible subdivision into rock layers. Using our techniques, such an illustration is created in a matter of minutes. It can be used to express an interpreter's internal idea of the layering. When the expressive illustration is made, his ideas are more easily grasped by others. More about such illustrations and Figure 11 in particular can be read in Paper V.

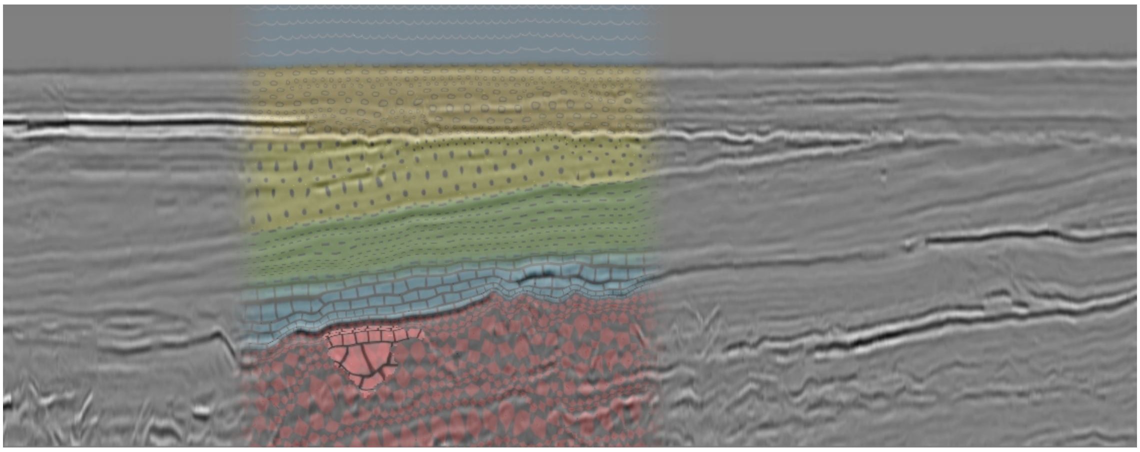

Figure 11: Seismic data overlaid with textures representing a hypothesized rock layering.

3.1.3 Improved Seismic Visualization

In the two next sections results are presented achieving improved seismic visualizations by adopting more realistic lighting models during volume rendering and results for 3D texturing.

Ambient Occlusion In Paper IV we introduce a new volumetric rendering method for seismic data based on ambient occlusion. The method is inspired by optical realism and achieves more natural shadowing, better depth perception of the data and allows for glowing effects to emphasize important regions. A major challenge for providing useful 3D interactive visualization is the choice of an appropriate 3D rendering algorithm. Seismic data are typically noisy and lack distinct material boundaries as the acquisition is based on acoustics. The widely used gradient-based methods as introduced by Levoy [12] are in general sensitive to high-frequency noise. Gradient-based shading of seismic data introduces distracting artifacts which makes interpreting the 3D renderings difficult. Other approaches, such as unshaded direct volume rendering or maximum intensity projection [15] tend to depict seismic data as a homogeneous cloud without distinct features. Thus, common approaches are frequently unsuitable for visualizing seismic data. Until now only unshaded direct volume rendering and gradient-based methods have been used for 3D seismic data. We have identified a volume rendering technique that is promising for seismic data which has not been able to be rendered in real-time until recently. The method is an approximation of ambient occlusion rendering. The approximation enables interactive rendering of the volume with real-time modifications of the transfer function. See Figure 12 for a comparison of gradient-based (left) and ambient occlusion based (right) rendering.

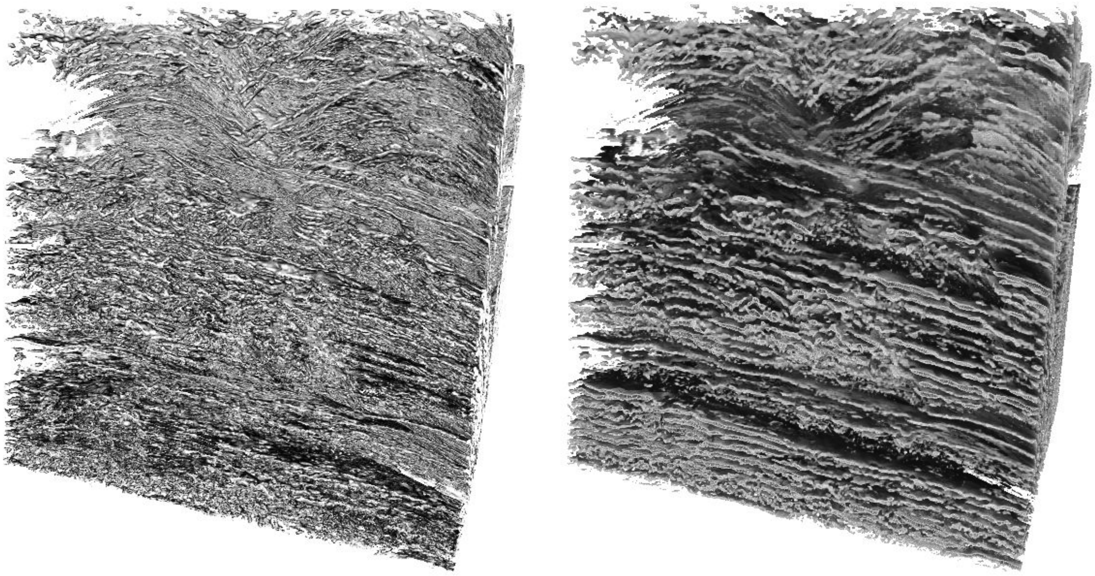

Figure 12: Standard gradient-based rendering to the left and our real-time approximation of gradient-free ambient occlusion to the right.

3D Sparse Texturing Paper V discusses the possibilities to extend the widely used 2D texturing in the geosciences to 3D by using sparse 3D textures. All standardized geologic textures are defined in 2D (Figure 4). However the seismic data to be interpreted is inherently 3D. 2D textures lend themselves directly to illustrate 2D seismic slices but have limitations when applied on 3D data. The current solution for 3D data is to apply 2D textures on planar cross sections. This technique is seen in 3D geological illustrations such as in Figure 3b). Several advantages can be gained if suitable 3D textural representations of the 2D textures are specified and used. From an algorithmic point of view it is simpler to map 3D textures to 3D volumes and surfaces than 2D textures. Distortion problems arise when mapping 2D textures to curved surfaces, and frame-to-frame incoherencies arise when interactively moving a 2D textured cut-plane. 3D textures will reduce these problems of spatial and frame-to-frame incoherencies. Using 2D textures on semitransparent volumetric regions is not well defined. Here, 3D semitransparent textures may give rise to a higher perceptual depth. Volumetric variations can be revealed inside the 3D texture and not only on the exterior boundary as is the case when using a 2D texture. However since the textures used for conveying knowledge in geology are only defined in 2D, there exist no corresponding 3D versions. If we want to use 3D textures, we must synthesize them ourselves from the 2D cases which is an underdefined task. Figure 13a) shows two 2D textures to the left and how they can be synthesized into 3D textures and used in a 3D rendering. For the limestone texture two alternative 3D textures are shown, a sparse texture at the top and a dense texture below. By also using parameterization information, the textures can be deformed and will communicate the deformation in 3D. Figure 13b) gives an example of a deformed 3D texture situated in a rock layer. The 3D texture exemplar is shown in the black rectangle in the lower left corner. The topic of 3D sparse texturing is discussed at the end of paper V.

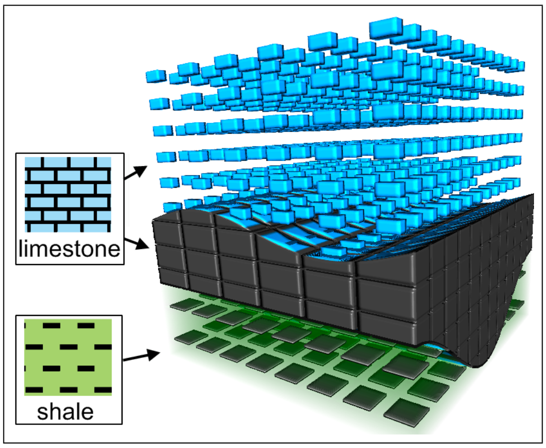

(a) Undeformed 3D textures synthesized from 2D textures.

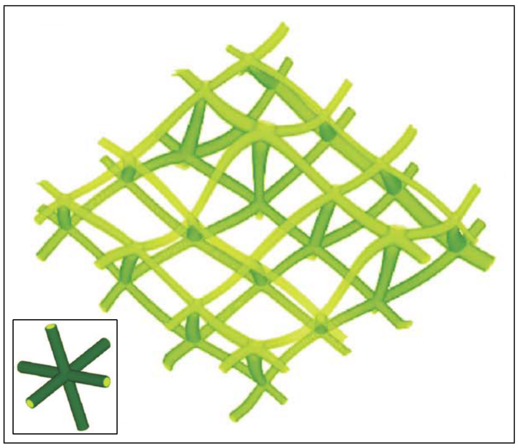

(b) Deformed 3D texture. The 3D texture exemplar is in the bottom left corner.

Figure 13: 3D sparse texturing allows seeing into the volumetric data.

3.2 Results in Rapid Interpretation

Building on the techniques reviewed in the previous sections and on automatic interpretation algorithms introduced in this section, methods for rapid seismic interpretation are presented here.

3.2.1 Overviews

Techniques for scale invariant and multiattribute visualizations result in overviews. Overviews can quickly identify where to focus the interpretation. Or they may early reveal if a seismic prospect has no potential and should be discarded. Either way, time is saved.

Scale Invariant Zooming Figure 14 gives two examples, to the left and right of the vertical line, of what we have called 'scale invariant' zooming. An image contains a certain amount of information (top row Figure 14). Information is averaged away or compressed beyond comprehension when simply zooming out an image. For efficient presentation, different zoom levels require different abstracted visualizations. We achieve this with our scale invariant zooming techniques. We first extract an abstraction of the information to be visualized, we use precalculated horizons and the parameterization for this. Then the information is presented at varying sparseness degrees by visualizing an appropriate amount of information for the current zoom level. Left shows normal zooming with the

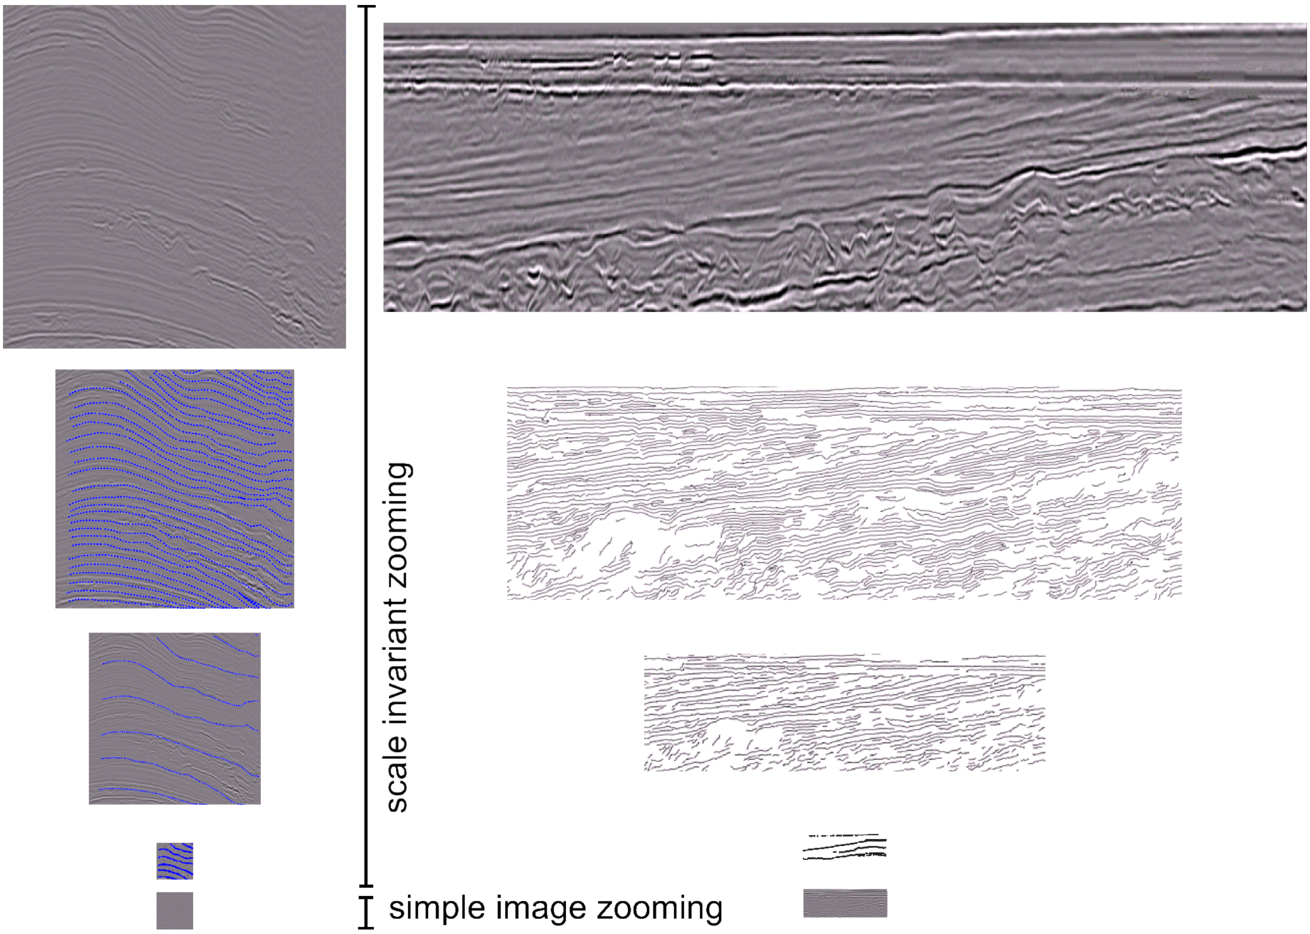

Figure 14: Examples of adaptive visualizations appropriate for different detail levels.

sparse information overlaid in blue while right shows only the sparse information. The bottom row shows the original data from the top row zoomed out in the normal way. In the bottom row literally all information is gone and only a gray rectangle is left.

Scale invariant zooming can also be used when zooming in on data. Using the method of texturing for instance, even in extreme sub-sample magnification, when zooming beyond the resolution of the data, textures will look smooth and will express the angle and the type of layer that is being zoomed in on. On such sub-resolution scales, color transfer functions would yield uninformative single colored results. The techniques for scale invariant zooming are described in Paper II.

Multiattribute Visualization Multiattribute visualization deals with visualizing several measurements in the same image. In many domains, such as medical visualization, climate visualization, or flow visualization, disparate measurements exist for the same area. In medical visualization one might have CT, MR, and PET volumes of the same organ. Being able to create a visualization that merges these volumes into one meaningful image is the topic of multiattribute visualization. For climate visualization, one has the challenge of displaying measurements such as temperature, wind direction, wind speed, and cloud factor on a geographic map already full of information. Extensive research has been done in multiattribute visualization. Bürger et al. [7] provided a state-of-the-art report. The challenge in multiattribute visualization is to merge all the information in one comprehendible image. The image can be made comprehendible to the viewer by mapping the attributes to representations that are perceptually separable from each other. More about perceptual orthogonality in textures and patterns can be read in Bertin [5]. In the seismic domain one must deal with data such as overlapping volumetric attributes, well logs and various formats that interpreted data is represented in, e.g., geometric surfaces.

Figure 15a) shows a slice of seismic reflection data. Derived from this data are the attributes chaos, dip, frequency and reflection intensity (Figure 15b - e). A multiattribute visualization of these attributes is shown in Figure 16. The reflection intensity attribute is depicted with a line transfer function. The line transfer function represents data by using curves following the parameterization which are colored according to the attribute value. Another line transfer function is used to draw stippled blue lines following the trend of the data. More can be read about line transfer functions in Paper II.

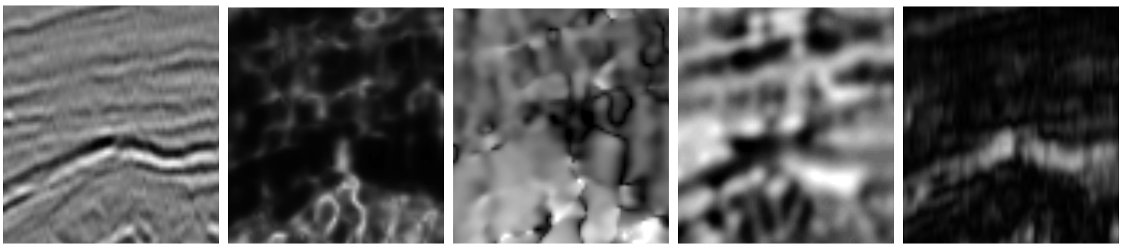

(a) reflection data (b) chaos (c) dip (d) frequency (e) reflection intensity

Figure 15: Reflection data in a) followed by attributes derived from it in b)-e).

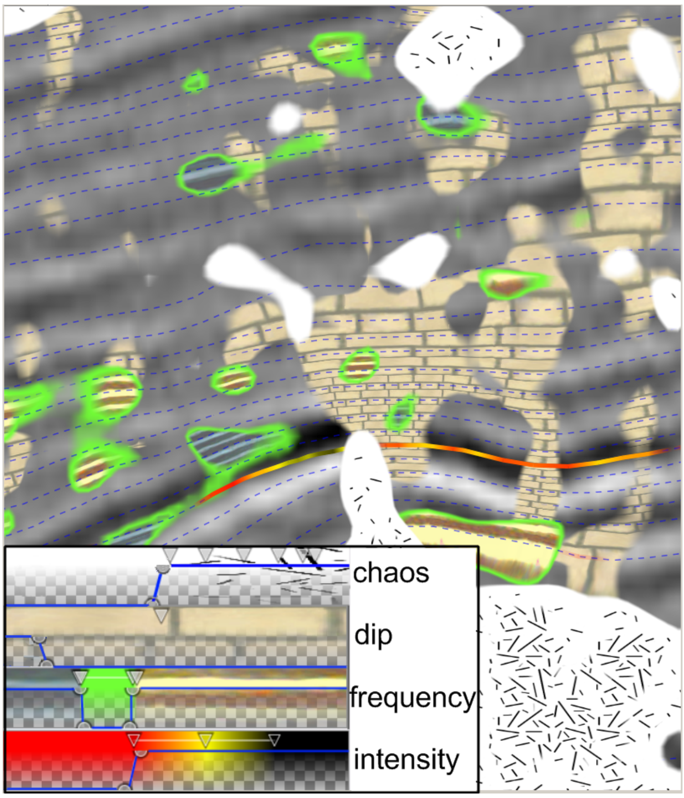

Figure 16: A multiattribute visualization with the attributes chaos, dip, frequency, and reflection intensity overlaid on seismic reflection data. The transfer functions are shown in the lower left corner. Notice the low and high frequency values with differently colored textures and a green halo effect easily obtained with the transfer function.

Figure 17 a)/d) is an example of using two differently colored versions of the same texture to encode intermediate (brown) and high (blue) dip values. High chaos values are overlaid semitransparently with an increasing amount of lines representing increasing chaos. Figure 17 b)/e) and c)/f) apply highly transparent and thus less disctracting textures. The transparency makes it possible to see the underlying reflection data. In example c)/f) small peaks in the transfer function opacities define the texture borders by making them more opaque so the interior regions can be made even more transparent.

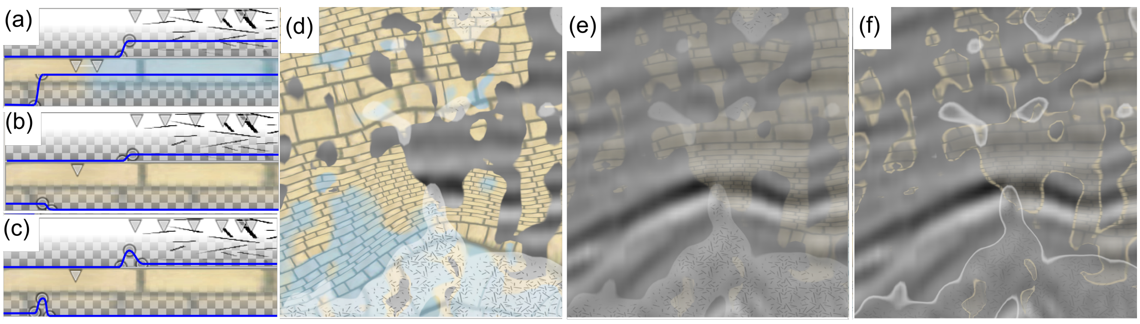

Figure 17: In a), b and c) three transfer function setups are shown each consisting of two transfer functions. The top transfer function maps chaos and the bottom one maps dip as given in Figure 15. Their corresponding renderings are shown in d), e) and f).

3.2.2 Automatic Structure Identification

The subsurface of the Earth consists of material layers with distinct densities and porosity characteristics. Horizons are central structures in interpretation defined as the interfaces between these layers. We have chosen to focus on seismic horizons since these are typically the first structures to be interpreted. They are also some of the simplest structures to identify using image processing techniques. This is due to their well defined visual appearance in the seismic data. With our methods we aim to increase the computer's assistance in finding horizon structures. Thereby enabling rapid interpretation.

Horizon Identification Computer-assisted interpretation of horizons is done by first automatically identifying horizon candidates in a preprocessing step and then presenting the candidates to the user through an intuitive user interface. The preprocessing step and the user interface for horizon selection is described in Paper II for the case of 2D slices. Preprocessing and interaction is extended to 3D in Paper IV.

By considering the reflection values of a 2D seismic slice (Figure 18a/b) as height values in a terrain (18c) one can identify the horizons as valleys and ridges (18d). We automatically trace out the valleys and ridges and create connected curves from the traces.

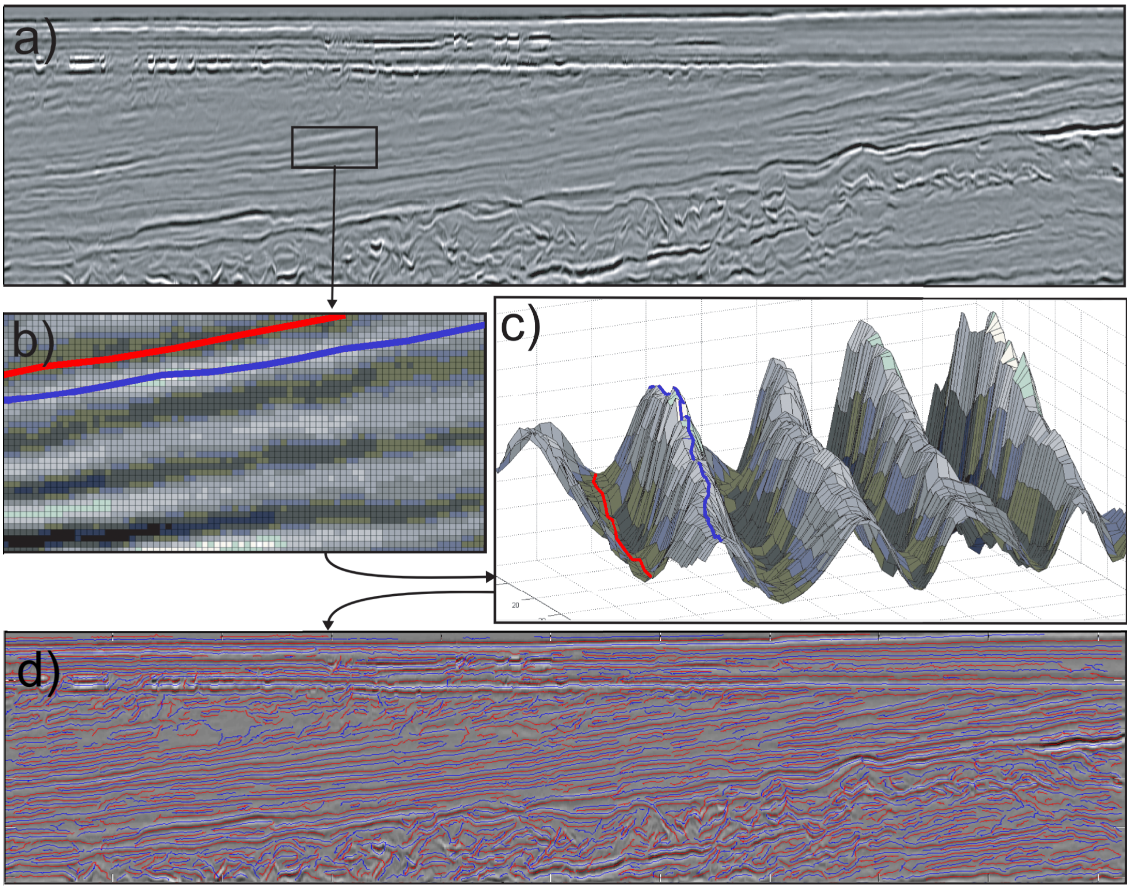

Figure 18: The 2D horizon extraction algorithm tracing along ridges and valleys in the reflection data). In c) a height field of the rectangle in a) is shown. A ridge is marked with red and a valley is marked in blue in b) and c). All extracted ridges in red and valleys in blue are overlaid on the reflection data from a) in d).

We extended the 2D horizon algorithm into 3D to enable rapid interpretation of 3D seismic horizons. A naive extension of the 2D method into 3D did not succeed. During horizon growing, too many unrelated horizons merged into single horizon candidates. After preprocessing a complete seismic volume, this could result in a single merged structure consisting of the union of all horizons found in the data. To resolve the problem, we performed a splitting of the produced horizon candidates using a hierarchical mesh-subdivision method [3]. There are many ways to split up a surface into smaller ones. Our subdivision was steered by maximizing the flatness of each surface part. After subdivision, the split horizon parts are selected in real-time by the geoscientist and assembled together into more appropriate horizons. The preprocessing steps consisting of horizon growing and subdivision are shown in Figure 19. For more information see Paper IV.

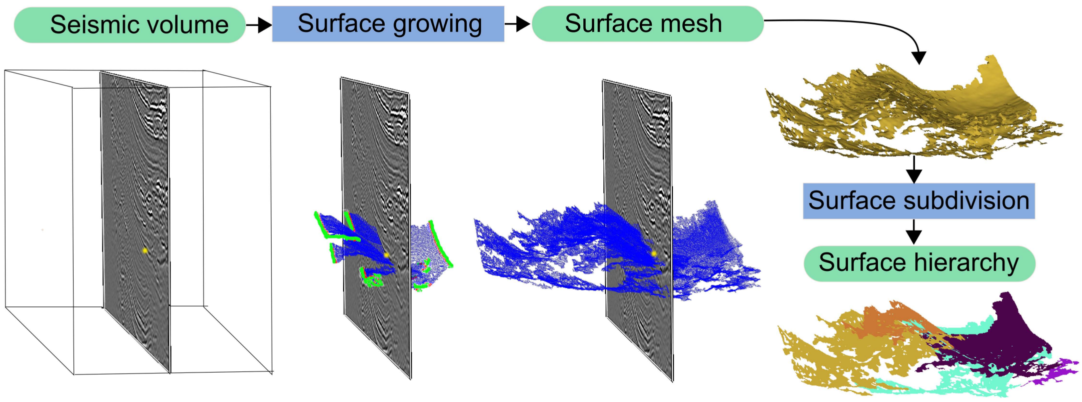

Figure 19: 3D horizon extraction followed by surface subdivision. Three steps of horizon growing from one seed point (yellow) is shown to the left. Bottom right shows the subdivision for one of the hierarchy levels.

Structure Identification by Similarity Search Instead of tailoring the structure identification algorithm for horizons only, a general method for extracting arbitrary structures based on similarity was attempted as inspired by the work of Borgos et al. [6]. Our method works by letting the user select a point on a seismic slice. Then the system shows all other points that have a similar vertical neighborhood to the selected point. The underlying assumption is that certain structures such as horizons, have distinct neighborhoods. Thus by selecting a point on a horizon, all other points on the horizon would be identified due to their distinct and similar neighborhood to the comparison point. The method works by evaluating each sample point based on its local properties only. Therefore this method is parallelizable as opposed to the sequential method of tracing horizons along ridges and valleys described in Figure 18. We were able to implement a version of the algorithm on the highly parallel GPUs of modern graphics cards. Thereby real-time performance is achieved and preprocessing is avoided.

The neighborhood that is compared is the n vertical samples above and below a specific location called comparison point. Thus n defines the extent of the neighborhood and is chosen by the user. A closeness score to the selected point is calculated for each sample in the seismic slice. This results in a new derived attribute which we call similarity attribute. A color transfer function is used on the similarity attribute and the result is overlaid on the original slice. Two examples are shown in Figure 20. The two examples have the same comparison point positioned on a horizon, but different color transfer functions. The comparison point is shown as a red dot. Its neighborhood values are shown as a yellow graph with sample values mapped to the horizontal axis. To verify that the horizon has a consistent neighborhood, the neighborhood graphs of four other points on the horizon shown in green are plotted in black. The extent of the neighborhood is shown as a colored vertical line through each point. The similarity or closeness metric has a large effect on the results. Our distance function is the sum of the squared distances between neighborhood components. We use this simple function that is quick to evaluate for achieving interactivity.

The method did not perform as well as we had expected. It was not possible to pinpoint the structures we selected with sufficient accuracy. Either a subset of the structure was selected (undersegmentation) or too much was selected (oversegmentation). Figure 20 a) shows a transfer function with transparencies set for maximizing the selection of the interesting horizon while minimizing the selection of other structures. The structure is undersegmented at the same time as areas not part of the structure are selected. Figure 20 b) shows a transfer function with more opacity defining a less strict similarity criterion. In this example we increased the transfer function opacity until the whole horizon was colored. But now oversegmentation is extensive. Thus capturing the horizon with this method is approximate only. The neighborhood along a horizon (using our metric) varies too much. At the same time due to the large amount and variation of other horizons, unrelated horizon fragments are matched. Figures 20 a) and b) show that the approach with vertical neighborhoods based on Euclidean distances is not able to identify horizons uniquely.

While trying to improve the method we noticed that the result is very sensitive when using other angles than a vertical neighborhood line. Using an angle that is normal to the horizon would better capture its neighborhood and could lead to a more consistent signature. Several heuristics with varying qualities exist for identifying the normal of a seismic horizon from the reflection data. Instead of using any of these, we extended the horizon signature to a version that was rotationally invariant by considering a spherical neighborhood instead of along a vertical line. The new method had similar problems as the

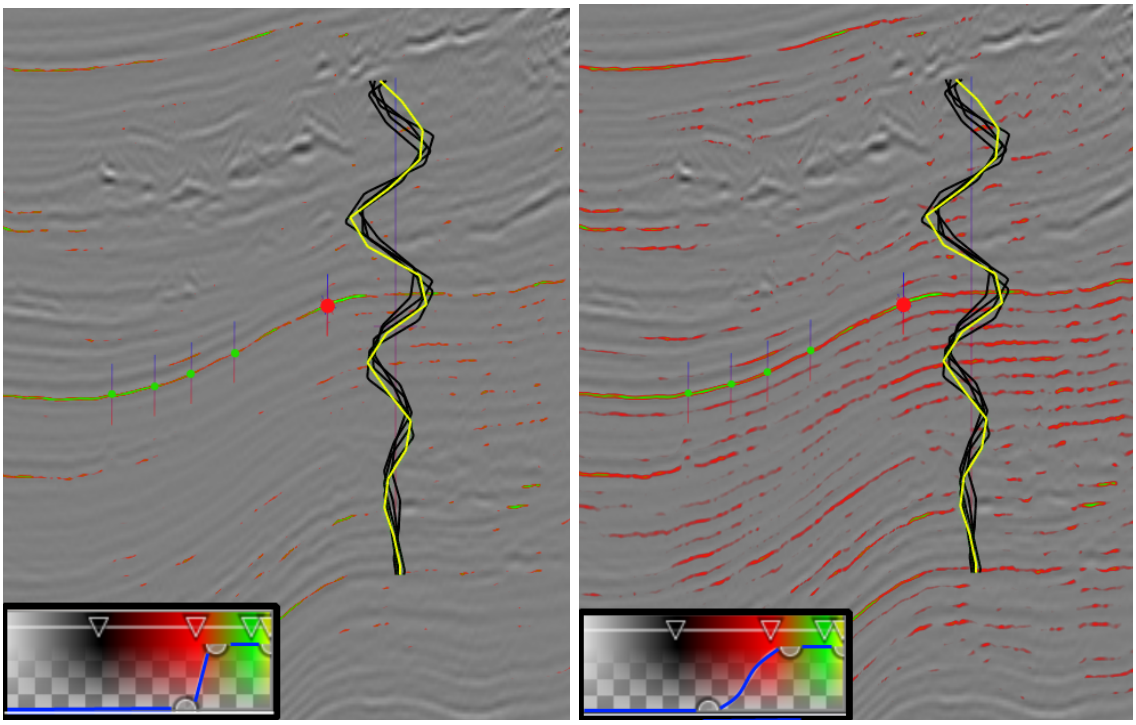

(a) (b)

Figure 20: Horizon extraction by neighborhood similarity. The horizon intersecting the red dot is attempted extracted. Its vertical neighborhood intensities are plotted in the yellow graph. Intensities of the green dots are plotted in black graphs. In a) a strict similarity criterion defined by the color transfer function in the lower left corner is used. In b) a less strict similarity criterion is used. Oversegmentation is visible in both examples.

previous one. However it proved itself useful on other modalities which will be discussed in the next section.

Moment Curves Although not directly related to seismic data, this work originated from the rotational invariant signature on seismic data just described. By considering the evolution of the mean and variance in a spherically growing neighborhood around the sample positions we were able to achieve promising segmentations of CT and MR data. Each voxel in a volume is assigned a sequence of the mean and variance values of the voxels in a spherical growing neighborhood. We then assign optical properties to the voxel based on this sequence of values. This results in a novel classification of materials and material boundaries. The evolution of the mean and variance as the spherical radius around a voxel increases, defines a curve in 3D for which we identify important trends and project it back to 2D along the radius axis. The resulting 2D projection can be brushed for easy and robust classification of materials and material borders. See Figure 21 for a 2D projection of the 3D curves. Figure 23 shows brushing and the corresponding classification on two coronal slices through a male CT dataset. The blue dots in Figure 21, the characteristic archs and why only the interior of the organs are classified in Figure 23 is explained in Paper III.

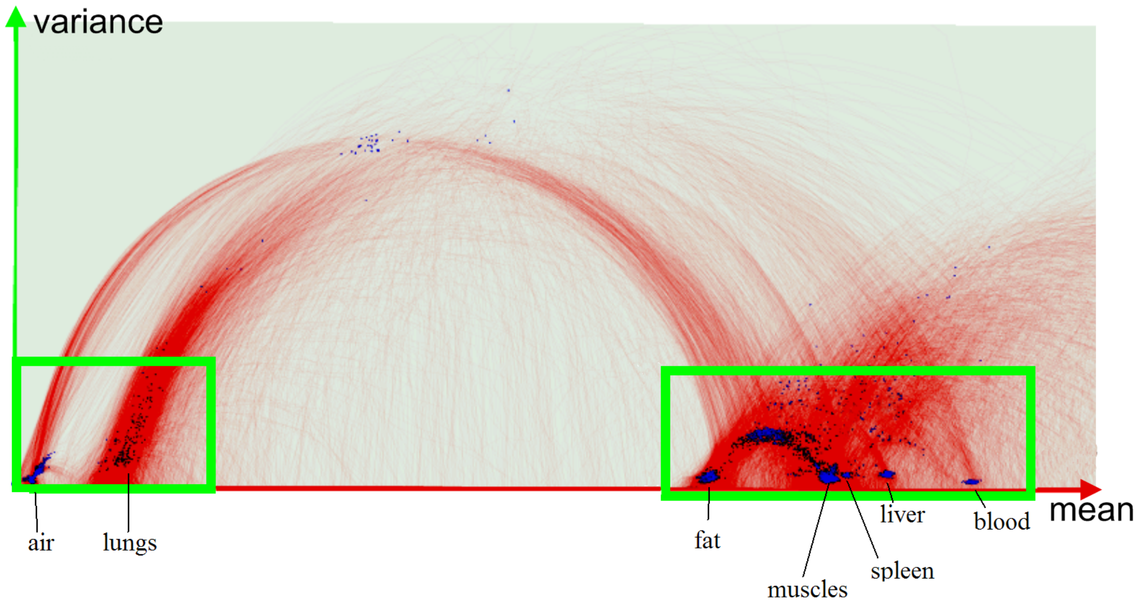

Figure 21: In red is seen (mean,variance,radius) curves projected into the mean/variance plane.

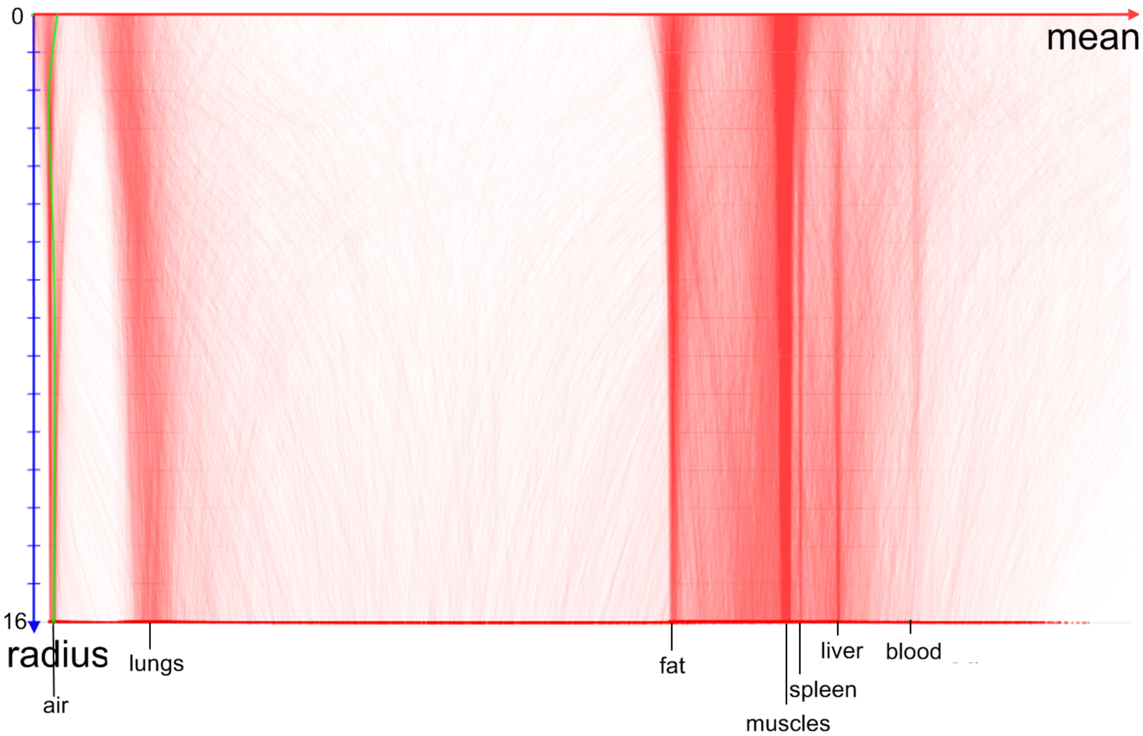

Figure 22: The (mean,variance,radius) curves projected into the mean/radius plane.

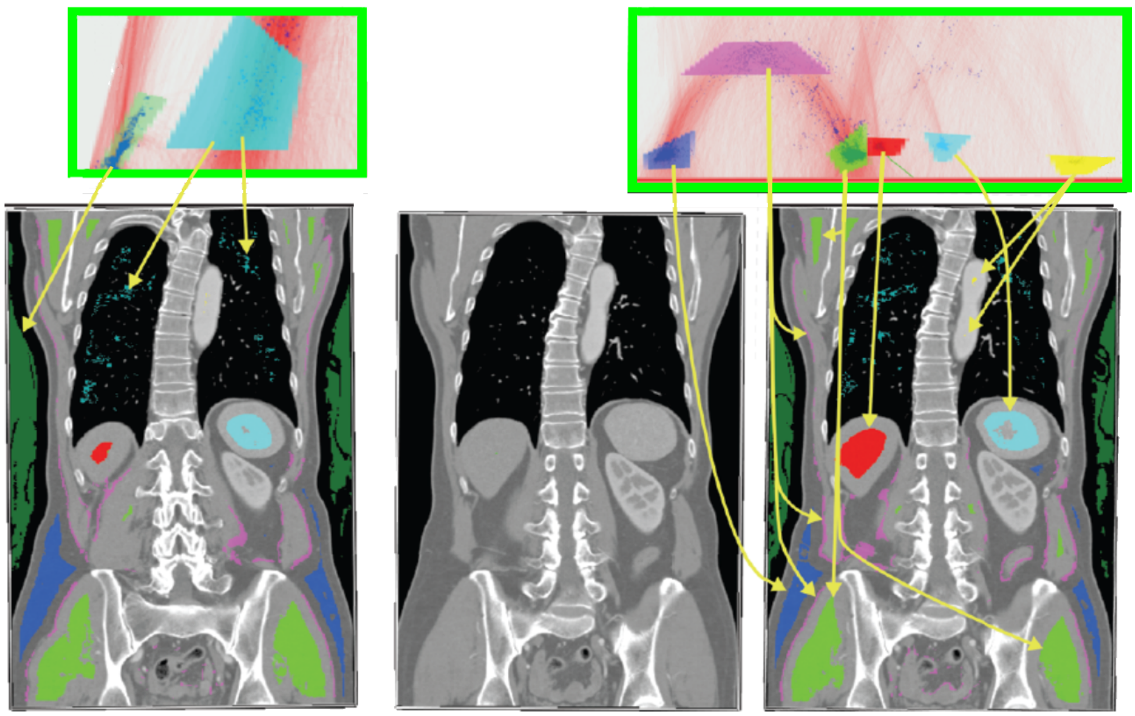

Figure 23: Brushing in the two indicated green regions from Figure 21 is shown. Bottom left and right shows two slightly different CT slices and their classified regions. Middle shows an unclassified version of the right slice.

3.2.3 Verification

Finally our methods can be used to verify a finished (Figure 24) or an ongoing interpretation hypothesis (Figure 25) with the underlying data. This is achieved by using transparency to smoothly move from the uninterpreted data to the interpreted expressive visualization as seen in Figure 24. The approach allows for comparison of the underlying data with the interpretation and thus enables verification. In Figure 25 a suggested layer subdivision has been overlaid the seismic data with high transparency. The layer borders are drawn opaquely.

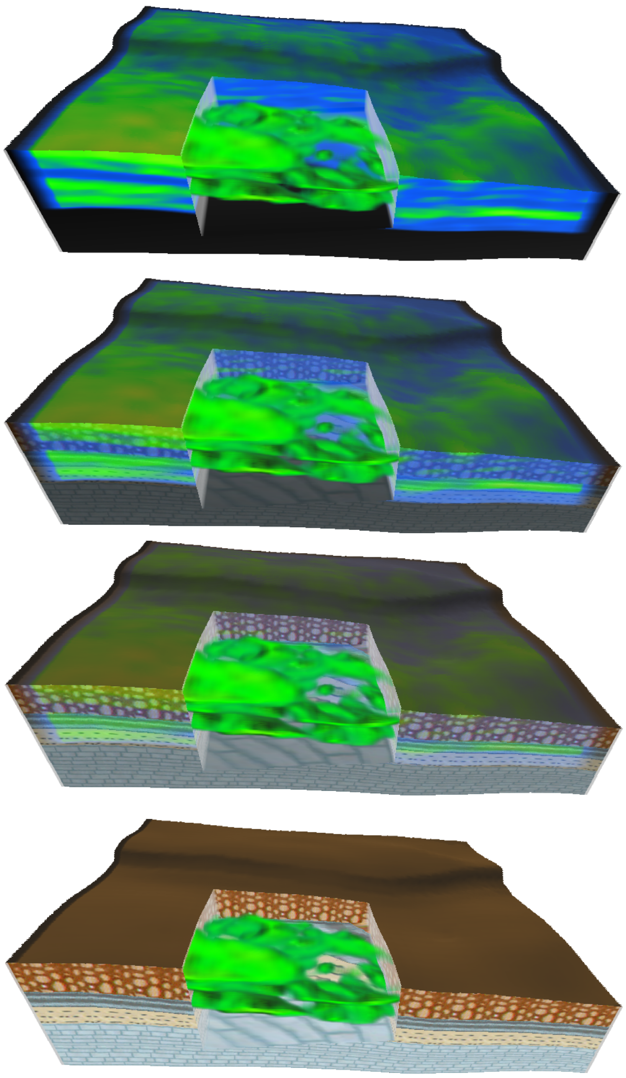

Figure 24: Verification of a finished interpretation. The top figure shows the 'impedance' attribute. Volume rendering of the attribute is shown in the cutout. At the bottom is shown an interpretation of layers.

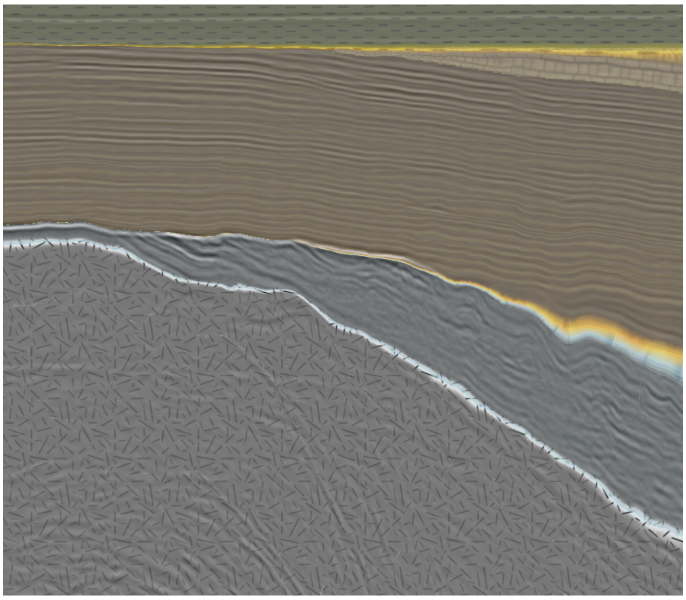

Figure 25: Verification of an ongoing interpretation.

Chapter 4

Conclusions

In this thesis we showed that geoscience has a well developed and established symbolic language for representing geo-information and that the use of manually crafted illustrations is an important part of the field. Methods for automatic generation of geoscientific illustrations were established. These methods can reduce the time to create illustrations which are used in reports and in education. The fact that computer-generated illustrations can be interacted with in real-time opens up for a more pedagogic presentation than with hand drawn static illustrations. Quickly generated illustrations also makes it possible to externalize hypotheses, making it easier for interpreters to communicate their hypotheses and inner models.

We indicated that rapid interpretation can be achieved using our expressive scale-invariance visualizations. Scale-invariant and multiattribute visualization techniques can provide overviews of the data. With such overviews one might be able to avoid interpretation in unimportant areas and spend more time interpreting areas with potential.

We presented methods for automatic structure identification of seismic horizons and, as a side track, of human tissue. Time spent on repetitive and time consuming tasks can be reduced with the automatic methods so the interpreter can focus on challenging areas in the data where automatic methods fail.

New ways of visualizing 3D seismic data were presented. Our real-time approximated ambient occlusion rendering gives enhanced visualizations of challenging high frequency and noisy seismic data. We presented the potential for 3D sparse texturing to convey rock layering, rock types, and deformations in true 3D. We also presented methods for verifying a final interpretation by seamlessly visualizing it with the underlying data the interpretation was based on.

The following paragraphs highlight observations made in this thesis.

The right level of automation A recurring theme of the thesis has been the strategy of identifying and automating work that computers can perform so humans may focus on what they do best. Here it has been important to strike the right balance by, on the one side not having an overautomatized process where the user will question and not understand the computer's suggestions, and on the other side not underautomatize the process and burden the user with time consuming monotonous tasks. We believe that the combination of computerized brute force preprocessing and a simple and responsive user interface, where the preprocessed proposals are presented to the user, helps achieve this goal. High-level preprocessing attempts to simulate human expertise. Perfect simulation of human expertise is impossible so errors are made. It is therefore important that the user can easily avoid selecting the wrongly generated suggestions with an efficient user interface.

The power of abstraction Another important point in the thesis is the use of preprocessed information that describes the structure of the seismic data. This information is of higher semantic level than derived seismic attributes. Derived seismic attributes show different properties of the seismic data but they do not give higher-level insight. For instance, in derived attributes there is no grouping of samples into objects. Higher-level information is required for further analysis and advanced visualization. We use the higher-level horizon and parameterization information for texturing, to create abstracted views of the data such as the different detail views in scale-invariant visualization and for rapid interpretation of horizons.

By randomly changing parameters and observing how this affects the visualization, the user in effect browses the reflection data and gets a better understanding of it. It is a common procedure to perform this browsing by altering view parameters and transparency parameters of the transfer function. However to get a structural overview, this real-time parameter modification has not been possible because of the need of manual intervention to extract structures such as horizons. Using the preprocessed information, the user can now quickly change parameters (such as the horizon transfer functions described in Paper II) that affect the visualization of structures and get a deeper overview .

The merge of analysis and report creation The thesis has focused on the workflow in geosciences consisting of the collection of 3D seismic reflection data, seismic interpretation, and visualization. Several concepts from this thesis can be adapted to other domains that also have these three phases, stated in more general terms, as data collection, analysis for finding results, followed by report making for communicating the results. Any domain with such stages might adapt methods presented in this thesis for the following advantages. During analysis, expressive sketching possibilities can help brainstorming and hypothesis making of scenarios. When expressive sketches made during analysis are of a quality comparable to illustrations found in reports, the tasks of analysis and report making are in effect merged. Thus the report is created continuously during analysis and not started from scratch when the analysis is finished. With expressive electronic sketches, the analysis stage is no longer a purely mental stage or a stage only semi-documented with rough pen and paper sketches. The analysis stage is better documented since the expressive visualizations represent the gained knowledge more explicit than rough sketches do.

The future - Interactive Reports Development in hardware is changing the physical appearance of reports. Computers, displays and interaction devices are merging and shrinking in size. In recent years we have seen the transition of computing power from stationary computers to laptops, netbooks and mobile phones. The size of the electronic devices reports are presented on are approaching the size of printed paper documents. We might be seeing the start of technology that will allow touch sensitive bendable electronic paper with integrated wireless network.

Development in software is changing the static nature of reports as demonstrated in this thesis. Currently, reports are typically static documents whether displayed on screens or on paper. There are increasingly examples of dynamic documents as interactive web content. A first step towards interactive reports with expressive visualizations of 3D data can be seen on the web page made by Jean-Paul Balabanian [4]. In his web page, instead of displaying a static image of a human skull, the image is generated in real-time from the underlying 3D CT scan of a head. The data can be interacted with by rotation and by changing a transfer function. Thus any aspect of the underlying head data can be visualized interactively directly in the document.

Using the new hardware and software technology together one can envision a change from static documents to dynamic and interactive hand held documents. This can lead to a new generation of reports with interactive illustrations. By integrating the underlying data with the analysis software into the report, all steps of the analysis can be gone through or even corrected. Thus analysis and report reading can be merged. We therefore conclude: In the future, data analysis and report creation will merge and all steps will be accessible and modifiable in the resulting physical document.

Bibliography

[1] *Commision for the geological map of the world.* http://ccgm.club.fr/index.html, 2006.

[2] *Federal Geographic Data Committee, Digital Cartographic Standard for Geological Map Symbolization.* Document FGDC-STD-013-2006, 2006.

[3] M. Attene, M. Spagnuolo, and B. Falcidieno. Hierarchical mesh segmentation based on fitting primitives. *The Visual Computer*, 22(3), sep. 2003.

[4] J.-P. Balabanian. Online realtime volume raycaster. `http://www.ii.uib.no/vis/team/balabanian/project/ORVR.cgi`, 2008.

[5] J. Bertin. *Semiology of graphics.* University of Wisconsin Press, 1983.

[6] H. G. Borgos, T. Skov, and L. Sønneland. Automated structural interpretation through classification of seismic horizons. In *Mathematical Methods and Modelling in Hydrocarbon Exploration and Production*, pages 89–106. Springer Berlin Heidelberg, 2005.

[7] R. Bürger and H. Hauser. Visualization of multi-variate scientific data. In *EuroGraphics 2007 State of the Art Reports (STARs)*, pages 117–134, 2007.

[8] J. D. Foley, A. van Dam, S. K. Feiner, J. F. Hughes, and R. L. Phillips. *Introduction to Computer Graphics.* Addison-Wesley Professional, 1993.

[9] J. Grotzinger, T. H. Jordan, F. Press, and R. Siever. *Understanding Earth.* W. H. Freeman and Company, 1994.

[10] M. Hadwiger, J. M. Kniss, C. Rezk-Salama, D. Weiskopf, and K. Engel. *Real-time Volume Graphics.* A K Peters, 2006.

[11] A. Iske and T. Randen, editors. *Atlas of 3D Seismic Attributes, Mathematics in Industry, Mathematical Methods and Modelling in Hydrocarbon Exploration and Production.* Springer, Berlin Heidelberg, 2006.

[12] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8:29–37, 1987.

[13] M. Livingstone. *Vision and Art: The Biology of Seeing.* Harry N. Abrams, 2002.

[14] E. R. Tufte. *The Visual Display of Quantitative Information.* Graphics Press, May 2001.

[15] J. W. Wallis, T. R. Miller, C. A. Lerner, and E. C. Kleerup. Three-dimensional display in nuclear medicine. *IEEE Transactions on Medical Imaging*, 8(4), 1989.